

By programming the knowledge bases and problem-solving techniques of human experts, researchers in artificial intelligence have created computer systems that may outperform the experts

# EXPERT SYSTEMS

## Programming Problem-Solving

ROBIN WEBSTER and LESLIE MINER

Preface	16
Business	20
Technology	22
Applications	24
Follow-up	28

In the North Sea an oil rig automatically shuts down in response to a failure somewhere in its maze of electrical and mechanical systems. Is the problem minor or serious? Time and money are at stake while the oil company waits to find out.

In the United States a geologist looks at the data from field studies of a particular site. Is it rich enough in minerals to warrant excavation?

Typically such problems would be handled by human experts, because they involve decision-making. But now these decisions can be tackled successfully by a new type of computer program called an expert system.

Expert systems are unlike conventional programs for number-crunching computers; they are the first systems designed to help humans solve complex problems in a commonsense way. They are being applied not only in the fields of oil exploration and geology but also in medicine, chemistry, genetic engineering and business, and they will eventually provide advice in the home. Soon you will be able to buy expert programs for your home computer that will help you to fill out tax forms, to plan investments or to diagnose (and perhaps repair) what is wrong with your car.

Until recently computers have been useful mainly in areas where human capabilities are limited. Large computers excel in high-speed

ROBIN WEBSTER and LESLIE MINER are free-lance science writers based in London.

calculating and in storing and retrieving enormous amounts of data. The computers run algorithmic programs—completely defined, step-by-step procedures for solving problems.

However, such exact computations can solve only a fraction of human problems. The challenge is to devise machines that are capable of "thinking," or at least of emulating the way humans employ past experience to solve new problems.

Expert systems are unique in that they can draw conclusions from a store of task-specific knowledge. They are called knowledge-based expert systems because they contain the same kind of rules used by human experts when they make decisions in their field of expertise. Expert systems encode knowledge in a symbolic form; they can be said to handle symbols. These systems draw conclusions principally through logical or plausible inference, not by calculation.

The most significant thing about an expert system is that it can display on request its line of reasoning in a human-language form. This quality is called "transparency," as opposed to the opaque, or black-box, quality of an algorithmic computer program.

Expert systems are not locked into any specific decision path; they can pick from alternative paths in their search for a conclusion. They can weigh facts and assumptions and in a sense make choices appropriate for each problem presented to them.

**WHAT ARE EXPERT SYSTEMS?**  
THEY ARE AN OUTGROWTH OF RESEARCH ON ARTIFICIAL INTELLIGENCE. EXPERT SYSTEMS ARE COMPUTER PROGRAMS WITH A VAST STORE OF SPECIALIZED KNOWLEDGE. THEY SOLVE COMPLEX PROBLEMS BY IMITATING THE DECISION-MAKING PROCESSES OF HUMAN EXPERTS. AND IN SOME INSTANCES THEY ALREADY OUTPERFORM HUMANS.

**WHY DO I WANT TO KNOW MORE ABOUT EXPERT SYSTEMS?**  
BECAUSE YOU WILL SOON BE ABLE TO USE EXPERT SYSTEMS IN YOUR BUSINESS TO SOLVE PROBLEMS. OR YOU MAY USE THE PROGRAMS ON YOUR HOME COMPUTER TO GET EXPERT ADVICE ON FILING TAX FORMS, OR CAR REPAIRS OR ON MEDICAL MATTERS.

**IS IT HARD TO LEARN TO USE AN EXPERT SYSTEM?**  
NOT AT ALL. BECAUSE THE SYSTEM WILL ASK YOU QUESTIONS IN PLAIN ENGLISH AND EXPLAIN ITS REASONING IF YOU ASK IT TO.

EXPERT  
SYSTEMS

64

Technology

**Codified Reasoning**

What has surprised the experts and computer scientists who have developed expert systems is how much of an expert's judgment and "gut feeling" can be codified. "If you ask a doctor or a geologist about his reasoning," says John G. Gaschnig, who heads an expert system project at SRI International, "he will tell you it is based on 'years of experience.' Now if we can capture those rules of thumb on paper along with his textbook knowledge, then we can get a computer to do the same thing."

Another role for expert systems will be in helping to explain and to service advanced computer systems and other machines that are becoming too complicated for most humans to work with directly. "By the end of this decade," says Thomas T. Kehler, manager of knowledge-based systems at Texas Instruments, "a computer system will typically have a million components. How can we possibly design these systems, let alone run and maintain them, without expert systems to help?"

**Transparent Reasoning**

Randall Davis of the Massachusetts Institute of Technology describes the need for transparent systems, which will display their line of reasoning on request: "What would you do if a program printed out an enormously long formula in answer to a question about how to build a bridge? Would you go ahead and build the bridge and then drive over it—would you accept it as a black-box result? You'd rather not, I'd rather not. I'd rather be able to ask the system, 'How do you know that? Where did that piece of information come from?'"

"Also, imagine the problems you have when a

standard software system fails. Instead of poring over long printouts or probing the machine's memory, you can go to the terminal and ask, 'What happened, why did you do that?' Imagine how much easier it is to debug a system that has some ability to backtrack, to recapitulate its action for you."

Expert systems are designed to display information from their knowledge base in human-language form and to indicate their line of reasoning in arriving at a conclusion. "The key thing about expert systems," says Donald Michie of the University of Edinburgh, "is that they not only provide accurate answers but also justify their answers in terms that make sense to the human user."

**Japan's KIP Computer**

Interest in expert systems is worldwide. Work is being done in Europe, the Soviet Union, Japan and the United States. Most of the advanced research is going on at a few major research centers in the United States. Especially noteworthy is Japan's Fifth-Generation Computer Development Project, in which a knowledge-information-processing (KIP) computer is being designed. The Japanese may design the computers with non-von Neumann architectures. KIP-based application systems are expected to handle many forms of input: speech, handwritten characters, video images and printed text. Most significant, the systems will be able to extract meaning from the information they are given, answer user questions and offer advice that may lead to solutions. The Japanese are said to be investing at least half a billion dollars in research and development for the fifth-generation machines.



Scottish computer scientist Donald Michie is a pioneer in expert systems.

ROBIN WEBSTER

## EXPERT SYSTEMS TECHNOLOGY

## How an Expert System Is Created

For decades artificial intelligence was the province of theoreticians and experimental programmers working almost unnoticed in the back rooms of academia. At first research in artificial intelligence sought a universal technique—a philosopher's stone—to conquer problem-solving. Called the weak-methods approach, it has to a large extent been abandoned.

Current research is directed at developing high-performance problem-solving programs for specific tasks. Formerly, AI had little

to offer industry. Now industry is suddenly fascinated with expert systems. As Peter E. Hart of Fairchild Camera and Instruments Corporation observes: "It has taken AI 25 years to become an overnight success."

A prerequisite for building and maintaining an expert system is an expert—a point that is often overlooked. The success of an expert system depends on whether it captures not only textbook knowledge but also what have come to be known as the heuristics of the human expert—intu-

ition, informal rules, experiential knowledge. All of this information makes up the knowledge base of the system.

The effectiveness of any expert system is actually determined by the quality of its knowledge base: its completeness, usefulness and validity. The task of codifying the expert's judgments is called knowledge engineering. A knowledge engineer must understand the expert's field to some extent, debrief him, translate his informal rules into IF-THEN rules and eventually formalize those rules

TECHNOLOGY/2

and his textbook knowledge into a computer program.

Knowledge engineering calls for a great deal of cooperation from the expert, and there is evidence that in turn his own expertise and understanding are enhanced by the process. A good example is the Prospector economic-geology program, which was developed at SRI International.

John G. Gaschnig, who heads the Prospector team, remarks that making an expert system is an exercise in the codification of a science. "Every geologist has commented on what a stimulating experience it was for him. We question the expert on all the relevant factors about his subject. We press him beyond where he is normally pressed by a client, and the process is often exhausting. One man said it made his head

hurt, but all agreed that it has made them understand their own subject better, made them better experts."

Another Prospector member, Rene Reboh of SRI, has developed a program called the knowledge-acquisition system (KAS). It is the means by which new knowledge is acquired and built into existing models or new models are created for Prospector.

Reboh describes it: "KAS makes it a lot easier to enter new models into Prospector, to change them around, to run them, to try out new things. When I first developed the program, I was pleasantly surprised. KAS knows the consequences of adding rules. It watches over you while you are building the system. It keeps tabs like a bookkeeper. You can ask if there is any unfinished business,

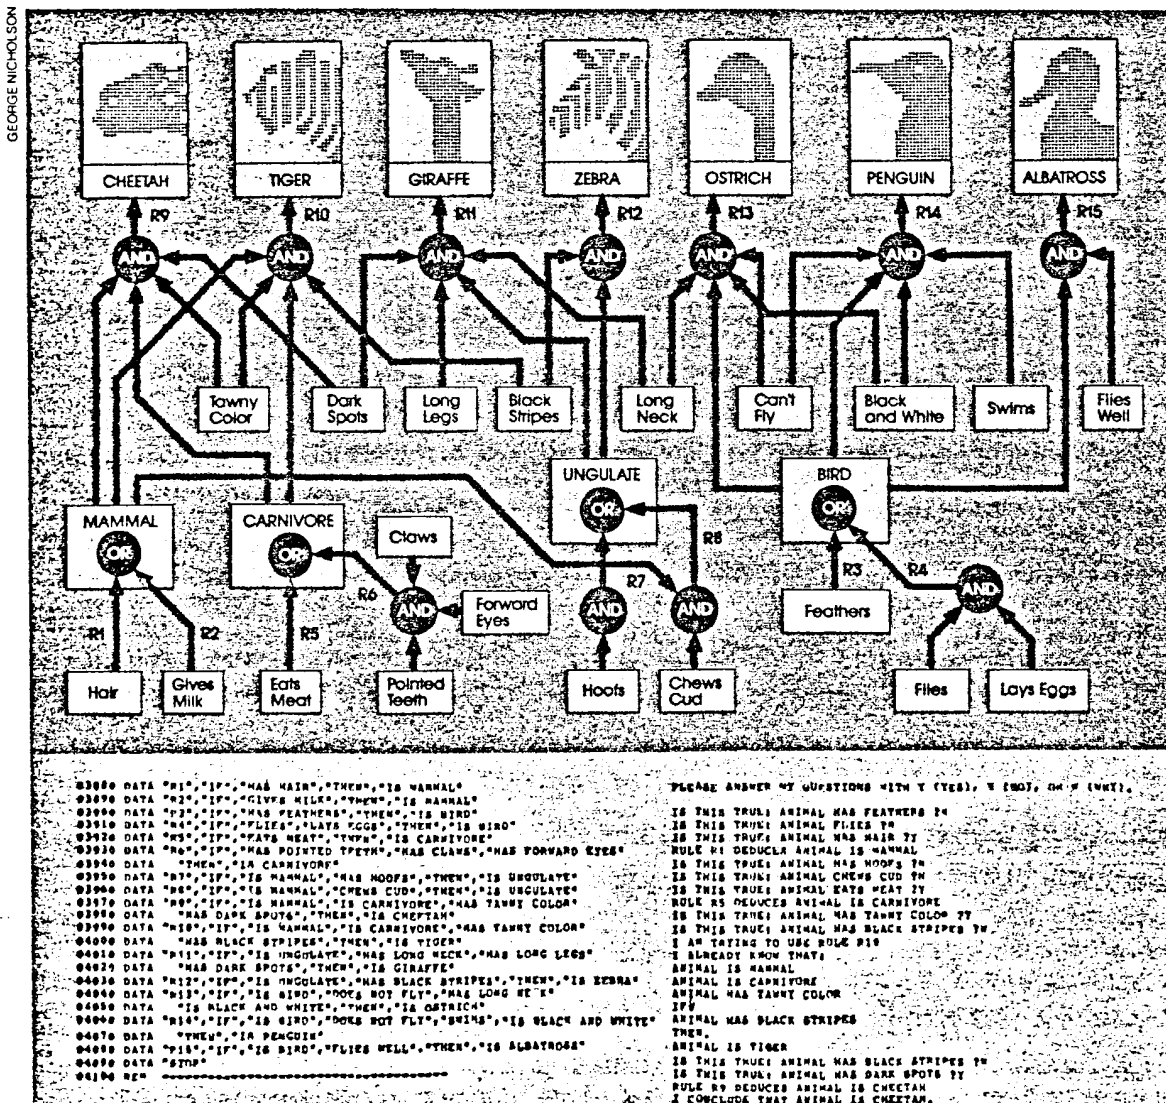
and it will tell you what information it needs. KAS could be adapted for use with other systems, but as of now it knows most about Prospector."

Other similar tools are the EMycin, RLL and Age systems in use at Stanford University, the Rosie system at Rand Corporation and the Expert system at Rutgers University.

To emulate human thinking a machine intelligence must be flexible and expandable. These are two hallmarks of expert systems. They can be expanded by the addition of new rules, and corrections to the knowledge base can be made and immediately assimilated. In this way an expert system continues to build up its knowledge base by adding new information incrementally. Thus the expert system is not static; it can be said to "learn."

## Is It an Albatross or a Zebra?

This inference network devised by Richard O. Duda and John G. Gaschnig is based on only 15 rules. The program identifies an unknown animal by starting with the hypothesis "animal is albatross" and then asking questions and making deductions.



EXPERT  
SYSTEMS

66

Business

## BUSINESS OUTLOOK

## First Profit-Oriented Expert Systems

One could imagine some use for expert systems in just about any sphere of business, engineering or research. Although it is still early to quantify the potential market for expert systems, industry is clearly discovering how valuable they can be in solving important problems. Take the case of Digital Equipment Corporation (DEC), a leading manufacturer of computer systems. According to John McDermott, a computer scientist at Carnegie-Mellon University, "The company has an a la carte business, and so each VAX-11 system ordered has a good chance of being unique." This marketing approach is profitable but troublesome. "Something like 20 percent of shipments had configuration problems, and to compensate for this the company tended to oversupply equipment." After three efforts using conventional computer methods, DEC tried an expert systems approach, almost as a last resort. It paid off.

Called R1, the system was developed by a group at Carnegie-Mellon, headed by McDermott. It is 90 percent accurate—more ac-



John McDermott heads a group that has developed expert systems for Digital Equipment Corporation.

curate than a human expert.

DEC now has a 15-man team figuring out where expert system techniques can be applied next. "In fact, we're helping DEC build a fairly major extension to R1 called XSEL," explains McDermott. "XSEL is a salesperson's assistant that will help our staff to identify the best type of DEC system for a customer. Future plans include developing a third system that will help manage the assembly and distribution of systems."

International Business Machines Corporation has set up ex-


pert systems projects to consider ways of making computer systems more reliable and easier to use. Several other big companies, including Xerox, Hewlett Packard and Texas Instruments, are working closely with research laboratories at universities such as MIT, Stanford and Carnegie-Mellon.

Schlumberger, the giant oil-field-services firm based in New York and Paris, offers a typical example of how knowledge-based computing has taken hold. "It's amazing how much expert geologists can tell you about what has occurred over millions of years at a particular site," says David R. Barstow, a computer scientist at Schlumberger. "This is an expert skill in oil exploration that needs to be captured in a system, and it was our first interest in AI."

Schlumberger is now applying expert systems in the development of software—for example, in solving the tricky problem of updating programs written in languages such as Fortran. "It is a question of obtaining knowledge about programming," says Barstow. "Can the expertise of human programmers be written down in rules that a system can use? It is important that we record the original design decisions. In our case, the changeability of software is a major issue. Most of it is based in empirically derived models and is subject to change all the time. As our understanding of the domain increases, the software has to change—we would build expert systems to facilitate that change."

As to the immediate payoff potential of expert systems, it should be noted that the Prospector team at SRI has made its first bona fide prediction—the location of an ore-body deposit in Washington State. John G. Gaschnig, head of the team, speculates: "It may be the first case of a mineral deposit discovered by computer." Events such as these will make awareness of expert systems widespread.

Programmed  
for  
Customer  
Service



Expert system R1 produces diagrams for technicians to follow in assembling each customized VAX-11 computer ordered from Digital Equipment Corp.

EXPERT  
SYSTEMS

67

Business

## What's What in Expert Systems

AREA OF APPLICATION	NAME OF SYSTEM	RESEARCHER(S)	DESCRIPTION OF SYSTEM
Computer Systems	R1	J. McDermott (Carnegie-Mellon)	Determines whether the set of components another system consists of is complete; determines relations among the components.
	XSEL	J. McDermott (Carnegie-Mellon)	An extension to R1. Assists salespeople in selecting appropriate computer systems.
Computing	PSI	C. Green (Kestrel Institute)	Automatic implementation of programs specified through natural-language dialogue.
	LIBRA	E. Kant (Carnegie-Mellon)	Efficiency-analysis component of PSI system.
	PECOS	D. R. Barstow (Schlumberger)	Knowledge-based program synthesizer for PSI system.
Education	GUIDON	W. J. Clancey (Stanford H.P.P.)	A computer-aided instruction program that teaches the student user by eliciting and correcting his answers to a series of technical questions.
Engineering	DART	(Joint project of Stanford H.P.P. and IBM)	Diagnostic-assistance reference tool. Used in field engineering. Investigates the use of expert systems to aid in the diagnosis of computer-system faults.
	KAS	R. Reboh (SRI)	Knowledge-acquisition system. Computer uses existing systems to design other systems.
	SACON	J. S. Bennett R. S. Engelmores (Stanford H.P.P.)	Structural analysis consultant. Assists structural engineers in identifying the best analysis strategy for each problem.
	SUX	H. P. Nii E. A. Feigenbaum (Stanford H.P.P.)	Identifies and evaluates moving objects (location, velocity, etc.) from primary-signal data.
Fault Diagnosis	AL/X	D. Michie (University of Edinburgh)	A form in which diagnostic experts can encode their knowledge of a specific domain, thus generating a system able to exercise knowledge on their behalf.
Geology	PROSPECTOR	P. E. Hart R. O. Duda (SRI)	A geological consultant that evaluates sites for potential mineral deposits. Based on CBC (computer-based consultant), a system that diagnoses electromechanical equipment.
	DIPMETER ADVISOR	(Schlumberger)	A system for analyzing information from oil-well logs based on production-system paradigms.
Knowledge Engineering	AGE	H. P. Nii N. Aiello (Stanford H.P.P.)	"Attempt to generalize." Guides the building of expert systems. Used in building PUFF.
	EMYCIN	W. Van Melle (Stanford H.P.P.)	A rule-based consultant that can be used in many fields. Derived from MYCIN.
	RLL	R. Greiner D. B. Lenat (Stanford H.P.P.)	Representation-language language. A tool for building expert systems quickly. An expert system whose expertise is knowledge representation.
	TEIRESIAS	R. Davis (dev. at Stanford H.P.P.)	Transfers knowledge from a human expert to a system and guides the acquisition of new inference rules.
Medicine	CADUCEUS	J. D. Myers H. E. Pople (University of Pittsburgh)	Diagnostic consultant program for different problems in internal medicine. Formerly known as INTERNIST.
	MYCIN	E. H. Shortliffe (Stanford H.P.P.)	Diagnoses bacterial infections and prescribes treatment.
	ONCOCIN	E. H. Shortliffe (Stanford H.P.P.)	Designed to assist physicians treating cancer patients. Guides management of complex drug regimens.
	PUFF	J. C. Kunz (Stanford H.P.P.)	Analyzes patient data to identify possible lung disorders.
	RX	R. L. Blum (Stanford H.P.P.)	Helps guide statistical analysis of chronic-disease patients whose data bases have been followed for a long time.
	VM	L. M. Fagan (Stanford H.P.P.)	Ventilator management. Uses physiologic transfers from patients to help guide the physician in determining when to take a patient off the ventilator.
Science	META-DENDRAL	B. G. Buchanan (Stanford H.P.P.)	Uses mass-spectrometry data to induce rules about the behavior of fragmented molecules.
	MOLGEN	J. Lederberg (Stanford H.P.P.)	Helps geneticists plan experiments involving structural analysis and synthesis of DNA.

Building  
Knowledge  
Bases

For a directory of research centers, see page 72.




## EMERGING APPLICATIONS

## Design, Diagnosis and Prospecting

The development of expert systems began at Stanford University in 1965, under the direction of Joshua Lederberg, Edward A. Feigenbaum and Bruce G. Buchanan. The first system, called Dendral, was not originally formulated as a knowledge-based system. The expert approach resulted from the failure of conventional computer methods in solving a formidable problem. The goal was to aid organic chemists in determining the molecular structure of compounds. Conventional methods had become so complex that the designers decided to try an approach now known as a knowledge-based system, using rules in the IF-THEN, or situation-action, form.

Today Dendral is used all over the world, especially by chemical companies. (Dow, Lederle and Merck were among the first businesses to show an active interest in expert systems.) Dendral has led to another program, called Meta-Dendral, which actually discovers new rules about the behavior of fragmented molecules. Meta-Dendral intrigued scientists because the induction of rules from data is basically a kind of theory formation. It particularly interested expert systems researchers, since any form of automatic rule-building reduces the time involved in knowledge engineering.

Another expert system, Mycin, diagnoses meningitis and other bacterial infections in the blood and prescribes treatment as well. Mycin was developed at the Stanford Medical School in the mid-1970's by



Dendral program generates possible structures for the chemical uvidin A; actual structure is No. 50.

Edward H. Shortliffe, a physician and computer scientist. It was a collaborative endeavor between the Stanford Medical School and the Heuristic Programming Project, also at Stanford. Meningitis is a disease particularly suited for expert system application. Quick diagnosis and treatment can save a patient from brain damage or death.

## Mycin: A Case Study


Mycin was originally developed to aid physicians during the critical 24- to 48-hour period when much of the decision-making is imprecise because all the relevant information is not yet at

hand. The program's record of correct diagnoses and prescribed treatments for meningitis and bacteremia has equaled the performance of human experts.

In a typical session the physician holds an interactive dialogue with Mycin about a particular patient (see box). He provides information and responds to questions generated by the program. The Mycin knowledge base comprises 500 rules in the IF-THEN form. A degree of certainty on a scale of 0 to 1 is supplied by experts to measure the importance of an association expressed in a rule. A user requesting a consultation can also apply a degree of certainty to qualify the answer. When the situation parts of the rules are satisfied, the action parts are "fired," or the conclusions of the rules are recorded.

Mycin examines evidence for and against approximately 100 diagnoses. The final output is a list of possible diagnoses above a suitable threshold of certainty. Mycin can take into account that adequate therapy must treat all likely infecting organisms and that at the same time it should not prescribe unnecessary drugs, such as antibiotics that could increase the possibility of toxic reactions or encourage the emergence of drug-resistant microbes.

Because physicians have the final responsibility for treatment, transparency in the system is essential. By typing RULE the user obtains a human-language explanation of the last rule executed. Typing WHY triggers a back-trace of the inference process that fired the system's last question. The number of the last rule is also given in case the user wishes to retrieve it. Repeated use of



Edward A. Feigenbaum, a founder of Stanford's expert systems project.

EXPERT  
SYSTEMS

69

Applications

1) \*\* NAME JOHN DOE AGE 34 SEX MALE RACE CAUCASIAN

2) ARE THERE ANY POSITIVE CULTURES OBTAINED FROM JOHN DOE?  
\*\* NO

3) ARE THERE ANY PENDING CULTURES OF JOHN DOE?  
\*\* YES

SITE DATE COLLECTED STAIN EXAMINED? METHOD OF COLLECTION (URINE, SPUTUM ONLY)

4) \*\* CSF 30 JAN 77 YES

STAIN OF SMEAR MORPH (TYPE NONE IF NO ORGANISM WAS SEEN ON THE SMEAR)

5) \*\* NONE (NO OTHER PENDING CULTURES.)

[CONSIDERING ORGANISMS (OTHER THAN THOSE ISOLATED FROM CULTURES) WHICH MIGHT BE PRESENT...]

13) HAS JOHN DOE RECENTLY HAD SYMPTOMS OF PERSISTENT HEADACHE OR OTHER ABNORMAL NEUROLOGIC SYMPTOMS (DIZZINESS, LETHARGY, ETC.)?  
\*\* YES

14) HAS JOHN DOE RECENTLY HAD OBJECTIVE EVIDENCE OF ABNORMAL NEUROLOGIC SIGNS (NUCHAL RIGIDITY, COMA, SEIZURES, ETC.) DOCUMENTED BY PHYSICIAN OBSERVATION OR EXAMINATION?  
\*\* YES

THE CSF CULTURES WILL BE CONSIDERED TO BE ASSOCIATED WITH MENINGITIS  
15) PLEASE GIVE THE DATE ON WHICH CLINICAL EVIDENCE (SYMPTOMS, SIGNS, OR LABORATORY TESTS) OF MENINGITIS FIRST APPEARED.

PLEASE ENTER CSF FINDINGS IN THE FOLLOWING TABLE

CSF WBC	PERCENT PMN	PROTEIN	CSF GLUCOSE	PERIPHERAL GLUCOSE
23) ** 100	99	280	5	UNKNOWN

34) IS MENINGITIS A HOSPITAL-ACQUIRED INFECTION?  
\*\* WHY

[2.0] THIS WILL AID IN DETERMINING THE ORGANISMS (OTHER THAN THOSE SEEN ON CULTURES OR SMEARS) WHICH MIGHT BE CAUSING THE INFECTION. IT HAS ALREADY BEEN ESTABLISHED THAT

- [2.1] THE INFECTION WHICH REQUIRES THERAPY IS MENINGITIS, AND
- [2.2] ORGANISMS WERE NOT SEEN ON THE STAIN OF THE PENDING CSF CULTURE, AND
- [2.3] THE TYPE OF THE INFECTION IS BACTERIAL, AND
- [2.4] NEUROSURGERY IS NOT ONE OF THE RELEVANT ITEMS FROM THE HISTORY OF JOHN DOE, AND
- [2.5] THE AGE OF JOHN DOE IS GREATER THAN 10 DAYS THEREFORE, IF

- [2.6] THE INFECTION WAS ACQUIRED WHILE THE PATIENT WAS HOSPITALIZED THEN THERE IS EVIDENCE THAT THE ORGANISMS (OTHER THAN THOSE SEEN ON CULTURES OR SMEARS) WHICH MIGHT BE CAUSING THE INFECTION ARE E. COLI ( 75) STAPHYLOCOCCUS-COAG-POS ( 3) PSEUDOMONAS-AERUGINOSA ( 3) KLEBSIELLA-PNEUMONIAE ( 5)

[BACK TO QUESTION 34...]  
\*\* YES

MY THERAPY RECOMMENDATIONS WILL BE DESIGNED TO TREAT FOR ORGANISMS THAT EITHER ARE VERY LIKELY OR, ALTHOUGH LESS LIKELY, WOULD HAVE A SIGNIFICANT EFFECT ON THERAPY SELECTION IF THEY WERE PRESENT. IT IS IMPORTANT TO COVER FOR THE FOLLOWING PROBABLE INFECTIONS(S) AND ASSOCIATED ORGANISM(S).

INFECTION-1 IS MENINGITIS

- \* <ITEM-1> E. COLI [FROM CLINICAL EVIDENCE ONLY]
- \* <ITEM-2> PSEUDOMONAS-AERUGINOSA [FROM CLINICAL EVIDENCE ONLY]
- \* <ITEM-3> KLEBSIELLA-PNEUMONIAE [FROM CLINICAL EVIDENCE ONLY]
- \* <ITEM-4> DIPLOCOCCUS-PNEUMONIAE [FROM CLINICAL EVIDENCE ONLY]

40) DOES JOHN DOE HAVE A CLINICALLY SIGNIFICANT ALLERGIC REACTION TO ANY ANTIMICROBIAL AGENT?  
\*\* NO

[CONSIDERING AMPICILLIN-AND-GENTAMICIN FOR USE AGAINST INFECTION-1...]

41) DO YOU HAVE REASON TO SUSPECT THAT JOHN DOE MAY HAVE IMPAIRED RENAL FUNCTION?  
\*\* NO

42) JOHN DOE'S WEIGHT IN KILOGRAMS (OR <NUMBER> POUNDS)  
\*\* 70

[REC-1] MY PREFERRED THERAPY RECOMMENDATION IS AS FOLLOWS:  
IN ORDER TO COVER FOR ITEMS <1 2 3 4>

GIVE THE FOLLOWING IN COMBINATION:

1) AMPICILLIN  
DOSE: 3.5 G (28.0 ML) Q4H IV [CALCULATED ON BASIS OF 50 MG/KG]

2) GENTAMICIN  
DOSE: 119 MG (3.0 ML, 80 MG/2 ML AMPULE) Q8H IV [CALCULATED ON BASIS OF 1.7 MG/KG] PLUS CONSIDER GIVING 5 MG Q24H INTRATHECAL  
COMMENTS: MONITOR SERUM CONCENTRATIONS

SINCE HIGH CONCENTRATIONS OF PENICILLINS CAN INACTIVATE AMINOGLYCOSIDES, DO NOT MIX THESE TWO ANTIBIOTICS IN THE SAME IV BOTTLE.

## Eaves-dropping on a Mycin Consultation

By typing a WHY question a physician consulting Mycin triggers the Teiresias program, which leads the user through the line of reasoning that produced Mycin's diagnosis.

REPRINTED WITH PERMISSION OF THE HEURISTIC PROGRAMMING PROJECT

EXPERT  
SYSTEMS

70

Applications

**Prospector:  
Accurate  
to within  
7 Percent of  
the Experts**

ISLAND COPPER  
DATA

STRUCTURES

BRECCIA PIPE




SULFIDE-ONLY  
STOCKWORKFAULT  
INTERSECTION  
REGION

FAULT

INTRUSIVE  
CONTACT

LIMIT OF SULFIDE

ISLAND COPPER

—FINAL  
FAVORABILITYOUTLINE OF  
ORE BODY

The expert system Prospector predicts the location of copper deposits based on early exploration data such as those shown in the geological map of a site in British Columbia (top). In response to Prospector's questions the user draws contours and fault lines on a digitizing tablet; the result is displayed on Prospector's screen (middle). Prospector then displays in orange or pink the areas that have the highest probability of containing ore. The outline of the actual copper deposit found after drilling is superposed on the prediction (bottom). Recently Prospector accurately predicted the location of a molybdenum deposit in the Mount Tolman area of Washington State.

these commands activates further back-tracing.

The importance of this capacity for explanation is described by Feigenbaum: "The intelligent agent viewpoint seems to us to demand that the agent be able to explain its activity; otherwise, the question arises of who is in control of the agent's activity. The issue is not academic or philosophical. It is an engineering issue that has arisen in medical and military applications of intelligent agents and will govern future acceptance of AI work in application areas."

Mycin, like Dendral, has already provided important tools for building future systems. One reason is that Mycin's proved mechanism for making inferences, called EMycin, can be used to create knowledge bases in other domains. The best example of this is Puff, an EMycin-driven lung-disease-diagnosis program that merely needed a set of rules about pulmonary functions in place of the Mycin knowledge base. Puff is a Mycin-like interactive system that ponders complex patient data in order to identify one of three lung disorders. Puff is in routine clinical use. Starting with about 55 rules, it was built in fewer than 50 hours of interaction with experts and 10 man weeks put in by the knowledge engineers—far less time than the usual two to 10 years required to build an expert system. The reason is that everything but the lung-disease knowledge of Puff is Mycin: the form of knowledge representation, the method of reasoning, the user interface, and so on.

A second offshoot of Mycin is Sacon, a basically EMycin-implemented system that uses rules from the field of structural engineering to provide advice about using a large engineering software system. Another is Guidon, a computer-aided instruction program that tutors students in a domain for which a program has been written for EMycin. Guidon has two separate knowledge bases: one consists of 200 rules about tutoring, the other of rules about a subject such as the field of infectious diseases. The system teaches the student user by eliciting and correcting his answers to a series of technical questions.

Also adaptable to other domains is Teiresias, a knowledge-acquisition program developed at Stanford by Randall Davis, who is now at MIT, to improve the Mycin knowledge base. When HOW and WHY are typed during consultations with Mycin, it is Teiresias that leads the physician user systematically through the line of reasoning that produced the computer's diagnosis. Conversely, Teiresias aids the physician in correcting or enhancing the knowledge base with new information, which the user simply types into the system. Teiresias maintains a record of Mycin's search tree for this purpose.

#### Prospector: A Case Study

The development of the Prospector expert system for geologists began in 1976 at SRI International in Menlo Park, California. The aim was to create a computer program that could serve as a consultant in the exploration for different types of ore deposits, such as porphyry copper, nickel sulfide and uranium.

John G. Gaschnig, who currently leads the Prospector project at SRI, says that the system offers a special challenge. "My interest is in assessing how

TECHNOLOGY/2




EXPERT  
SYSTEMS

71

Applications

## Training in Expert Systems for Sale



Bruce G. Buchanan of Teknowledge explains company's systems.

To meet the growing interest in AI, particularly in expert systems, the venturesome capitalists of California's Silicon Valley have begun to launch new companies that specialize in hard information about the new field as well as in the required software and design aids. The entrepreneurs include Edward A. Feigenbaum of Stanford's computer science department, who has joined his colleague Bruce G. Buchanan in starting Teknowledge, Inc., a company that sells not only sophisticated software but also training for industry in how to apply expert systems.

Many companies have already placed employees in Teknowledge training courses, which cost \$2,250 for one week and range up to \$38,000 with

a \$10,000 computer fee for a six-month course. Lee Hecht, president of the company, expects client employees who have computer programming experience to be able to return to their companies after six months and begin to build expert systems. After six months of instruction they will be able to use tools for building expert systems based on the EMycin and Age programs.

To promote the new company, and to help meet the growing need in industry for expert systems, Teknowledge is also offering one-day business-executive seminars throughout the United States. For more information write Dina Barr, Teknowledge, Inc., 151 University Ave., Palo Alto, CA 94301.

expert an expert system is," says Gaschnig. "How are human experts measured? If you play chess, you've got a number or ranking, but if you're a geologist or doctor, there aren't any commonly used yardsticks. We are taking several approaches with Prospector. In one we measure a particular ore model against the geologist who designed it—and if we can get the system to emulate the expert, we are happy." In a number of test cases geologists completed questionnaires about the value of a target area; the same data were fed to Prospector, which evaluated the certainty of the data on a scale of -5 to +5.

"We then make a detailed comparison between the human expert and the system," says Gaschnig. "We need to know not only if they agree but if they agree for the same reasons." Called knowledge debugging, this technique enables researchers to build an expert system incrementally, to make it more expert with time. "On the average," Gaschnig reports, "Prospector is accurate to within 7 percent of the experts, which I would say is pretty good."

One major advantage Prospector may offer over human experts is cost. A first-rate geological consultant might charge hundreds of dollars per day—and is well worth it, considering the potential payoff for the company that hires him. Yet at current commercial computer rates a typical consultation with Prospector costs only \$10 to \$15. When expert system programs become widely available for microcomputers, the price of a consultation will fall.

At this early stage, when programs such as Mycin and Prospector are still being developed, it is not clear what the cost of producing expert systems commercially will be. One simple measure of the business possibilities is the cost of hiring human experts to solve urgent problems. If it costs much less to create an expert system and maintain it than it does to consult an expert in person—as it most likely will—then the business potential is there.

All in all, the proved accuracy and expected economy of expert systems make it attractive to consider where they will appear next:

EXPERT  
SYSTEMS

72


Follow-up

Accuracy  
and Economy  
Determine  
Business  
Potential

► **Cancer chemotherapy.** A team headed by Shortliffe at the Stanford Medical School has developed an expert system called Oncocin. Designed to assist physicians who treat cancer patients, the system guides the management of complex drug regimens by monitoring the patient's response to treatments and prescribing the most effective chemotherapy series with the fewest toxic reactions.

► **Object identification.** Feigenbaum and H. Penny Nii of Stanford did a classified project for the Federal Government's Advanced Research Projects Agency (ARPA) involving the creation of an expert system, dubbed SU/X, that will identify and track certain moving objects based on the sketchiest of sensor information. The system, Feigenbaum and Nii say, "forms and updates hypotheses about location, velocity and other factors for an object in three-dimensional space from primary signal data."

► **Air-traffic control.** The control of airplane maneuvers is so complex and variable that algorithmic



Air-traffic-controller expert system is capable of planning the approaches of several landing airplanes.

computers are limited to providing information rather than control. A project at the University of Illinois is aimed at creating an expert system model of air-traffic-control decision-making.

► **Chemical development.** Companies such as Lederle need faster ways to determine and synthesize chemical structures. Through expert systems they can employ the ingenuity of the best chemical designers without having to mix a single vial.

► **Intelligent factories.** Mark S. Fox at Carnegie-

Mellon University is working with Westinghouse Electric Corporation's Robotics Technology Division on the idea of an "intelligent factory" that uses expert system techniques to ensure that all the machines are working at a high level of efficiency. The overall system would monitor scheduling of machine tasks to minimize downtime and configure the systems in such a way that work would flow through the plant as quickly as possible. The system will also monitor breakdowns and present the manager with the best alternative methods of production until repairs are made.

Expert systems will be applied, experts predict, to an incredible array of human problems. But just as algorithmic systems have their limits, so too will expert systems. Thomas T. Kehler of Texas Instruments puts the situation in perspective:

"Caution needs to be taken that industry does not expect too much out of AI. Having said that, an interesting transition takes place when you decide that you are interested in applying computers to help people to solve problems."

Douglas B. Lenat at Stanford is also cautious about AI, and yet he is intrigued by its promise. "It should be kept in mind that business people depend a lot on personal impact, personal communication, nuances. I don't see that being replaced. Nevertheless, AI will gradually be brought into the day-to-day running of things. Communication will be improved by developments such as automatic programming, where it will be possible to give short, stylized commands to a system, and thereby have it make dramatic changes to its program. Another aspect, which can be seen as complementary to automatic programming, is the modeling of individuals in programs."

Lenat speculates: "I look forward to having a semi-intelligent computer monitoring the house. It won't be a case of having lots of robots in the house, it will be more like the house as robot."

"Alan Newell of Carnegie-Mellon wrote a wonderful story a few years ago about the Land of Fairy. What he basically said was that the real excitement, the real promise of AI, is that all those fairy tales that we've heard about, talking lampposts and such, are not just fairy tales. Soon they will all start coming true." □

## FOLLOW-UP

## RESEARCH CENTERS

**Bolt Beranek and Newman, Inc.**  
10 Moulton St.  
Cambridge, MA 02238  
(617) 497-3361  
Contact: William A. Woods

**Carnegie-Mellon University**  
Department of Computer Science  
Schenley Park  
Pittsburgh, PA 15213  
(412) 578-2259  
Contact: John McDermott

**Digital Equipment Corporation**  
(DEC)

129 Parker St.  
PK 3-2 M18  
Maynard, MA 01754  
(617) 493-2775  
Contact: Peter D. Connell

**Fairchild Camera and Instrument Corporation**  
464 Ellis St.  
Mountain View, CA 94042  
(415) 962-5011  
Contact: Peter E. Hart

**Hewlett Packard**  
3000 Hanover St.  
Palo Alto, CA 94304  
(415) 857-5069

**IntelliGenetics, Inc.**  
124 University Ave.  
Palo Alto, CA 94301  
(415) 328-4870  
Contact: Martin N. Hollander

**Kestrel Institute**  
1801 Page Mill Rd.  
Palo Alto, CA 94304  
(415) 494-2233  
Contact: Cordell Green

**Machine Intelligence Corporation**  
330 Potrero Ave.  
Sunnyvale, CA 94086  
(408) 737-7960  
Contact: Earl D. Sacherdoti

TECHNOLOGY/2

EXPERT  
SYSTEMS

73

Follow-up

**Massachusetts Institute of Technology**  
Artificial Intelligence Laboratory  
545 Technology Square  
Cambridge, MA 02139  
(617) 253-5879  
Contact: Randall Davis

**Ohio State University**  
Department of Computer  
Information and Science  
2036 Neil Avenue Mall  
Columbus, OH 43210  
(614) 422-0923  
Contact: B. Chandrasekaran

**Rand Corporation**  
1700 Main St.  
Santa Monica, CA 90406  
(213) 393-0411

**Rutgers University**  
Department of Computer Science  
Hill Centre  
New Brunswick, NJ 08903  
(201) 932-2001  
Contact: Saul Amarel

**Schlumberger-Doll Research**  
Old Quarry Rd.  
Ridgefield, CT 06877  
(203) 438-2631

**SRI International**  
333 Ravenswood Ave.  
Menlo Park, CA 94025  
(415) 859-5494  
Contact: John G. Gaschnig

**Stanford University**  
Heuristic Programming Project  
Department of Computer Science  
Stanford, CA 94305  
(415) 497-2273  
Contact: Bruce G. Buchanan

**Teknowledge, Inc.**  
151 University Ave.  
Palo Alto, CA 94301  
(415) 326-6827  
Contact: Lee Hecht

**Texas Instruments**  
P.O. Box 225474  
Dallas, TX 75265  
(214) 995-2011

**University of California at Santa Cruz**  
Computer and Information Sciences  
237 Applied Sciences  
Santa Cruz, CA 95064  
(408) 429-2565

**University of Edinburgh**  
Machine Intelligence Research Unit  
Hope Park Square  
Edinburgh, Scotland EH8 9NW  
United Kingdom  
Contact: Donald Michie

**University of Illinois**  
Coordinated Science Laboratory  
1101 W. Springfield Ave.  
Urbana, IL 61801  
(217) 333-6071  
Contact: David L. Waltz

**University of Pittsburgh**  
Department of Computer Science  
Pittsburgh, PA 15260  
(412) 624-6475

**Xerox PARC**  
3333 Coyote Hill Rd.

JANUARY/FEBRUARY 1982

Palo Alto, CA 94304  
(415) 494-4000

**Yale University**  
Department of Computer Science  
10 Hillhouse Ave.  
328 Dunham Lab  
New Haven, CT 06520  
(203) 436-8160  
Contact: Steven Slade

## PERIODICALS

**AI Magazine**  
American Association for Artificial  
Intelligence  
445 Burgess Dr.  
Menlo Park, CA 94025  
(415) 328-3123  
Treatment of work in progress at the  
AI laboratories in the United States.  
Quarterly publication for members,  
\$10 per year.

**Sigart Newsletter**  
(Special Interest Group AI)  
Association for Computing  
Machinery  
P.O. Box 12115 Church St. Station  
New York, NY 10249  
(212) 265-6300

## BOOKS AND REPORTS

**Artificial Intelligence.**  
Patrick H. Winston. Addison-Wesley,  
Reading, MA, 1977.

**Artificial Intelligence and Natural  
Man.**  
Margaret Boden. Basic Books, New  
York, 1977.

**Expert Systems in the  
Microelectronic Age.**  
Donald Michie (ed.). Columbia  
University Press, New York, 1980.  
Gives a good overview of expert  
systems techniques that have been  
developed to date and the in-depth  
processes involved.

**The Handbook of Artificial  
Intelligence, Vol. I.**  
Avron Barr and Edward A.  
Feigenbaum (eds.).  
William Kaufmann, Inc., Los Altos,  
CA, 1981.

**Knowledge-Based Systems in AI.**  
Douglas B. Lenat and Randall Davis.  
McGraw-Hill, New York, 1981.

**Machines Who Think.**  
Pamela McCorduck.  
W. H. Freeman and Company,  
San Francisco, 1979.

**Proceedings of the International Joint  
Conferences on AI**  
American Association of Artificial  
Intelligence  
445 Burgess Dr.  
Menlo Park, CA 94025  
(415) 328-3123

**Prospector System Handbooks**  
SRI International  
333 Ravenswood Ave.  
Menlo Park, CA 94025  
Project reports on the development  
and operational facilities of the  
Prospector expert system.

## TECHNOLOGY RESOURCES

*Technology* magazine offers extensive  
amplification of this article and many  
others through our single-source in-  
formation service, *Technology Re-*  
*sources*. To order any of the following  
items, please turn to page 90.

(02301) **AI Comes of Age.** Edward K.  
Yasaki in *Datamation*, p. 48; Oct.,  
1980.

(02302) **AI for Systems Management.**  
Frederick Hayes-Roth. Rand paper se-  
ries P-6573, Rand Corporation, 1700  
Main St., Santa Monica, CA 90406;  
Jan., 1981.

(02303) **Annotated Bibliography  
about Prospector.** John G. Gaschnig.  
SRI International, 333 Ravenswood  
Ave., Menlo Park, CA 94025; Feb.,  
1981.

(02304) **Artificial Intelligence Labora-  
tory brochure.** Patrick H. Winston,  
Director, Artificial Intelligence Labo-  
ratory, MIT, Cambridge, MA; 1981.

(02305) **Conclusions from the Com-  
modity Expert Project.** James L.  
Stansfield. Artificial Intelligence Labo-  
ratory Memo 601, MIT, Cambridge,  
MA; Nov., 1980.

(02306) **Expert Systems.** Donald  
Michie in *The Computer Journal*, Vol.  
23, No. 4, pp. 369-376; Nov., 1980.

(02307) **Expert Systems: Newest  
Brainchild of Computer Science.** Ann  
M. Shoben. Rand paper series No. P-  
6552, Workshop on Expert Systems,  
San Diego, CA; Aug., 1980.

(02308) **Heuristic Programming Pro-  
ject brochure.** HPP, Department of  
Computer Science, Stanford Universi-  
ty, Stanford, CA 94305; 1980.

(02309) **R1: A Rule-Based Configurer  
of Computer Systems.** John McDermott.  
Paper No. CMU-CS-80-119, De-  
partment of Computer Science, Car-  
negie-Mellon University, Pittsburgh,  
PA 15213; Apr., 1980.

(02310) **R1—The Formative Years.**  
John McDermott in *AI Magazine*, pp.  
21-29; Summer, 1981.

(02361) **Basic Information Pack:** Ex-  
pert Systems. Reprints of the 10 ar-  
ticles and reports cited above plus a list  
of 20 additional citations. Price: \$45.

(02362) **Comprehensive Current Li-  
brary:** Expert Systems. Reprints of 30  
articles and research reports. Provides  
a detailed examination of artificial in-  
telligence and expert systems that are  
currently under development in uni-  
versities, research groups and corpora-  
tions. Includes material on expert sys-  
tem programming theory and case  
studies of expert systems now in use.  
Price: \$195.

(02301-02360) **Custom Reprint Ser-  
vice:** Expert Systems. Individual re-  
prints of the articles cited above are  
available on request. Please order by  
code number and article title. Price:  
\$8 per article; \$16 minimum order.

For More  
Resources  
Turn to  
Page 90